

Towards Quantifying Complexity with Quantum Mechanics

Ryan Tan,^{1,2} Daniel R. Terno,³ Jayne Thompson,² Vlatko Vedral,^{4,2,5} and Mile Gu^{1,2}

¹*Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, 100084 Beijing, China*

²*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543 Singapore, Singapore*

³*Department of Physics and Astronomy, Macquarie University, Sydney, New South Wales 2109, Australia*

⁴*Department of Physics, University of Oxford, Clarendon Laboratory, Oxford, OX1 3PU, United Kingdom*

⁵*Department of Physics, National University of Singapore, 2 Science Drive 3, 117551 Singapore, Singapore*

While we have intuitive notions of structure and complexity, the formalization of this intuition is non-trivial. The statistical complexity is a popular candidate. It is based on the idea that the complexity of a process can be quantified by the complexity of its simplest mathematical model - the model that requires the least past information for optimal future prediction. Here we review how such models, known as ϵ -machines can be further simplified through quantum logic, and explore the resulting consequences for understanding complexity. In particular, we propose a new measure of complexity based on quantum ϵ -machines. We apply this to a simple system undergoing constant thermalization. The resulting quantum measure of complexity aligns more closely with our intuition of how complexity should behave.

PACS numbers: 03.65.-w, 03.67.-a.

Are there any universal laws governing the evolution of complexity? While the second law of thermodynamics indicates ever increasing entropy, complexity seems to behave differently. The hot, smooth plasma near the Universe's birth and the final state of thermal equilibrium predicted by its heat death both appeal to our intuition of simplicity. Yet between these extremes, where there are stars, galaxies and life, the universe is complex; and because of that it is interesting. To answer this question, one must first quantify complexity.

This task is non-trivial. While we have intuitive notions of what is interesting or complex, they are deceptively difficult to formalize. The first attempts to quantify complexity came in the form of Kolmogorov complexity [1]. This measure equates the complexity of a sequence of numbers to the size of the minimal computer program that generates the sequence. While Kolmogorov complexity correctly identifies a constant sequence, consisting entirely of 0s, as simple; it is maximized by sequences that are completely random. This makes the measure unsatisfying for characterizing complexity, because it seems to misconstrue randomness with structure [2].

A promising way to avoid this problem came from the study of computational mechanics, where one is concerned in building ϵ -machines, the simplest predictive model of a supplied stochastic process [3]. One reason is that if a stochastic process is more complex, then replicating its future statistics will require more information about its past. A completely random process, for example, requires no past information to reproduce its future statistics; and nor would a process that outputs only zero. Meanwhile, a process with less trivial behavior, such as one which alternates between zero and one on successive emissions, can only be faithfully replicated by storing its last emission - and is thus more complex. The minimum amount of past information required to optimally predict

a given process thus introduces a more suitable measure of complexity. Known as the *statistical complexity*, C_μ , its clear operational significance and relative ease of evaluation have resulted in its widespread adoption in diverse contexts [4–9].

Nevertheless, statistical complexity still displays certain incongruities. Notably, it is not continuous – infinitesimal perturbations in the statistics of a process can lead to large jumps in its statistical complexity. A physical process that asymptotically approaches total randomness can have monotonically increasing statistical complexity, even if its final steady state has statistical complexity zero [6]. This seems to contradict our intuition of what complexity should be.

The results discussed so far however, have been limited to classical logic. Reality is ultimately quantum mechanical. If quantum logic allows us as to build simpler predictive models - models that use less past information - then the quantum analogue of statistical complexity could be a more fitting quantifier of structure. Indeed, it has been recently demonstrated that quantum models can optimally predict the future statistics of a classical stochastic process while generically storing less information about the past than their classical counterparts [10]. Systems that are complex to predict classically may be simpler quantum mechanically.

The objective of this article is to explore how quantum logic can improve our understanding of what makes a process interesting and complex, and ultimately contribute to discovering how complexity evolves. We review ϵ -machines - the provably simplest classical models; and how they can be further simplified through quantum logic. This motivates us to introduce a new measure of complexity, C_q , based on the complexity of quantum ϵ -machines.

We apply these ideas to a toy system featuring mono-

tonically increasing entropy. First we show that the system's complexity, when characterized by C_q , aligns more closely with our intuition of how complexity should behave. In addition to being zero when the system features zero or maximal entropy, it is also continuous; an infinitesimal perturbation in the statistics of the process will cause an infinitesimal change in C_q . Thus, we highlight the relevance of quantum mechanics in studying how structure and complexity persist and evolve.

The article is structured as follows. Section I will review statistical complexity, ϵ -machines, and their extension to quantum ϵ -machines. Section II reviews a new measure of complexity based on the complexity of these quantum ϵ -machines [10]. Subsection IIIB then applies this measure to a toy system with monotonically increasing noise; and highlights how the complexity of quantum and classical ϵ -machines diverge. Concluding remarks are presented in Section IV.

I. PRELIMINARIES

This section provides a background mechanism of inferring the statistical complexity of observed phenomena and their quantum extensions. For a more extensive treatment of these topics, see Refs. [4, 10, 11]. Familiarity with quantum information to the level of [12] is assumed.

A. Computational Mechanics, Complexity and Predictive Models

Computational mechanics seeks to study the complexity of systems through the lens of predictive models. The general approach is to assume that a system's behavior can be probed at discrete points in time t , with outcomes $x_t \in \Sigma$ dictated by random variable \mathbf{X}_t . Here, Σ defines the set of possible observable outcomes.

In the ideal scenario, the system may be probed indefinitely. The observable behaviour of such a system is a sequence of output values $\vec{x} = \cdots x_{-2}x_{-1}x_0x_1x_2\cdots$, where $\overleftarrow{x} = \cdots x_{-2}x_{-1}$ and $\vec{x} = x_0x_1x_2\cdots$ are the output sequences of the past and future respectively. This results in a stochastic process, defined by the joint probability distribution $P(\overleftarrow{X}, \vec{X})$. Here, \overleftarrow{X} and \vec{X} are the random variables governing \overleftarrow{x} and \vec{x} . Each realization of the system has a particular past \overleftarrow{x} with probability $P(\overleftarrow{X} = \overleftarrow{x})$ and exhibits a particular future \vec{x} with probability $P(\vec{X} = \vec{x} | \overleftarrow{X} = \overleftarrow{x})$.

Computational mechanics aims to infer the complexity of the system through these statistics. It asks, if we are to build a mathematical model of the process with statistically indistinguishable behavior what is the minimal amount of information it needs to keep about past observations? The more memory required the greater its complexity.

To formalize this intuition, envision our system encased in a black box, that simply outputs the outcome $x_t \in \Sigma$ at time $t \in \mathbb{Z}$. In a second black box, a computer attempts to simulate the process through execution of an appropriate mathematical model. It takes as input some S that is a function of past observations, \overleftarrow{x} and outputs appropriate future statistics. For this simulation to be completely faithful the two boxes must be indistinguishable. For each instance of the process with past \overleftarrow{x} , the model must output statistical predictions \vec{x} , according to the statistical distribution

$$P(\vec{X} | \overleftarrow{X} = \overleftarrow{x}). \quad (1)$$

The amount of information such a model requires to track is then determined by the minimal amount of space it needs to store about \overleftarrow{x} . Formally this is given by the information entropy of \mathbf{S} , the random variable that governs input variable S . We refer to this as the complexity of a given predictive model.

B. Classical Statistical Complexity

The statistical complexity for a process is determined by the complexity of its simplest model – that correctly simulates the process while requiring the least amount of information about \overleftarrow{x} . Thus we must find the best way of compressing the past without losing information about the future.

An immediate brute-force attempt is to store the entire past. Such a model takes the input \overleftarrow{x} directly and outputs the future according to (1) and thus stores the input with information content $C = H(\overleftarrow{X})$, where H denotes the Shannon entropy. This is clearly not efficient. For example, for a series of fair coin tosses $P(\overleftarrow{X}, \vec{X})$ is uniform over the distribution of binary strings. Thus C is infinite, implying that a simulation using this approach will require an infinite amount of memory. Clearly better approaches exist.

The simplest classical predictive models are epsilon machines (ϵ -machines) [3, 4]. Jointly proposed by Crutchfield and Shalizi, ϵ -machines are based on the reasoning that two different pasts need only be distinguished if they have differing future statistics. This motivates an equivalence relation on the set of pasts, \sim , such that for any two distinct pasts, \overleftarrow{x} and \overleftarrow{x}' , $\overleftarrow{x} \sim \overleftarrow{x}' \iff P(\vec{X} | \overleftarrow{X} = \overleftarrow{x}) = P(\vec{X} | \overleftarrow{X} = \overleftarrow{x}')$. Each equivalence class is referred to as a causal state S_i governed by a random variable \mathbf{S} . Let $\mathcal{S} = \{S_i\}_{i=1,2,\dots,N}$ denote the set of all causal states where N represents the total number of causal states. An ϵ -machine does not store \overleftarrow{x} , but only stores the equivalence class to which \overleftarrow{x} belongs. More formally:

Definition I.1 (ϵ -machines) *Given a stochastic process $P(\overleftarrow{X}, \vec{X})$, we can define its ϵ -machine as follows: The ϵ -machine of the process is the ordered pair $\{\epsilon, \mathbf{T}\}$,*

where ϵ is the causal state function such that $\epsilon(\overleftarrow{x}) = S_i$, $S_i \in \mathcal{S}$, $\mathbf{T} = \{T_{j,k}^{(r)} | r \in \Sigma, S_j, S_k \in \mathcal{S}\}$ is a collection of transition probabilities, with $T_{j,k}^{(r)} = P(S_t = S_k, \mathbf{X}_t = r | \mathbf{S}_{t-1} = S_j)$.

After initiating an ϵ -machine in state $\epsilon(\overleftarrow{x})$, there exists standard algorithms to systematically generate desired future statistics. At each time-step t , an ϵ -machine in causal state S_j will emit output $r \in \Sigma$ and transit to causal state S_k , with probability $T_{j,k}^{(r)}$.

Since there are no simpler classical models, Crutchfield and Shalizi defined the statistical complexity of a given process to be synonymous with the internal entropy of the ϵ -machine, i.e.,

$$C_\mu = H(\mathbf{S}) = - \sum_{i=1}^N p_i \log_2 p_i \quad (2)$$

where $p_i = P(\mathbf{S} = S_i = \epsilon(\overleftarrow{x}))$ is the probability that $\overleftarrow{x} \in S_i$.

Statistical complexity does not equate randomness with structure. For a completely random process the conditional future for all pasts is the same. Thus only one causal state is needed to specify the process and $C_\mu = 0$. This implies that completely random processes have no inherent structure. On the other hand, a process that emits a constant bit-string also features zero complexity as any probability distribution of a constant variable has zero entropy. In between these extremes, statistical complexity can be expected to peak.

Other properties of statistical complexity are more puzzling. Using a simple example of a general two-state process depicted in Fig. 1, for $q_0 = q_1 = 0.5 + \delta$, where $\delta > 0$ is arbitrarily small, the output is very close to being completely random. Nevertheless, the conditional futures of the two causal states differ, and $C_\mu = 1$. However, at $\delta = 0$, C_μ is 0. Such discontinuity appears surprising for a measure of structure.

C. Quantum is Simpler

All the above results assume our mathematical model processes classical information. Can a quantum extension of statistical complexity have different qualitative behavior?

Quantum logic allows the input information to be encoded into a quantum mechanical system. The advantage here is that this gives extra freedom to encode information in quantum superpositions. This lends a quantum refinement to the standard ϵ -machine. In quantum ϵ -machines each causal state is associated with a quantum causal state:

$$|S_j\rangle = \sum_{k=1}^{|\mathcal{S}|} \sum_r \sqrt{T_{j,k}^{(r)}} |r\rangle |k\rangle, \quad (3)$$

where $|r\rangle$ belongs to a Hilbert space of dimension $|\Sigma|$ and $|k\rangle$ to a space of dimension $|\mathcal{S}|$. The quantum causal states, $|S_i\rangle$, are in general not orthogonal; nevertheless we can construct a systematic method to sample from $P(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x})$, when given the appropriate quantum causal state $|S_i = \epsilon(\overleftarrow{x})\rangle$.

To see this, consider a machine that takes $|S_i = \epsilon(\overleftarrow{x})\rangle$ directly as its input. It can generate the the correct output statistics at each time step, by measuring $|S_j\rangle$ in the $|r\rangle$ basis. Each specific outcome r occurs with probability $T_{j,k}^{(r)}$, and collapses the system to $|k\rangle$. The machine sets the output x_0 equal to “ r ”, and the quantum ϵ -machine then applies a quantum operation that maps $|k\rangle$ to state $|S_k\rangle$. Iterating this protocol will give a series of output values $x_0 x_1 x_2 \dots$ with the correct statistical distribution. This can be formalized as:

Definition I.2 *The quantum ϵ -machine of a process $P(\overrightarrow{X}, \overleftarrow{X})$, is the ordered pair $\{\epsilon_q, \mathcal{S}_q\}$ where $\mathcal{S}_q = \{|S_i\rangle\}_{i=1,2,\dots,N}$ is the set of quantum causal states, and ϵ_q is a function such that $\epsilon_q(\overleftarrow{x}) = |S_i\rangle$.*

The complexity of the resulting model is again determined by the entropy of its input that is given by the von Neumann entropy

$$C_q = - \text{Tr } \rho \log \rho \quad (4)$$

where $\rho = \sum_i p_i |S_i\rangle \langle S_i|$ and p_i is the probability of initiating the machine in state $|S_i\rangle$. Since the $|S_i\rangle$ are generally non-orthogonal, C_q is often strictly less than C_μ . Quantum ϵ -machines thus can have complexity below what is possible using any classical predictive model.

The advantage of quantum ϵ -machines over their classical counterparts rests in the observation that future predictions do not require complete knowledge of the causal state the process started in. Indeed in the case of classical ϵ -machines two instances of a process can start in different causal states, S_i and S_j , where both $T_{i,k}^r$ and $T_{j,k}^r$ are non-zero. With probability $T_{i,k}^r T_{j,k}^r$ both ϵ -machines will transition to some coinciding causal state S_k upon the same emission r . If this happens, all future statistics will be identical and no amount of future observations can ever fully identify whether the system started in S_i or S_j . Thus some of the information used to distinguish between S_i and S_j is wasted.

Quantum mechanics allows the freedom to store different causal states as non-orthogonal quantum states, without employing classical randomness. Quantum ϵ -machines exploit this – they utilize quantum causal states that distinguish past causal states only to a degree sufficient for generating correct statistical behaviour. For example, for a process with two classical causal states S_0 and S_1 , the emission alphabet $\Sigma = \{0, 1\}$ and transition probabilities $T_{j,k}^r = \delta_k^r T_{jk}$ (where δ_k^r is the Kronecker delta) we have quantum causal states

$$|S_0\rangle = \sqrt{T_{0,0}^{(0)}} |00\rangle + \sqrt{T_{0,1}^{(1)}} |11\rangle \quad (5)$$

$$|S_1\rangle = \sqrt{T_{1,0}^{(0)}} |00\rangle + \sqrt{T_{1,1}^{(1)}} |11\rangle \quad (6)$$




FIG. 1: The causal state diagram for a two-causal state process. The causal states, S_0 and S_1 , are represented by a pair of circles. Arrows denote possible transitions. An arrow pointing from S_j to S_k , $j, k \in \{0, 1\}$, is labeled by the probability that the epsilon machine will transition from causal state S_j to causal state S_k upon emitting $x_t = r$: $T_{j,k}^r$, we also explicitly specify the matching value of r . For example, the above process has $P(S_0, \mathbf{X}_t = 0 | S_1) = T_{1,0}^0 = q_1$ with accompanying emission $r = 0$.

Since we are dealing effectively with a two-dimensional space we label $|00\rangle$ as $|\widetilde{0}\rangle$ and $|11\rangle$ as $|\widetilde{1}\rangle$. If S_1 and S_2 have non-zero probability of transitioning to the same causal state, S_k , then clearly $T_{0,k}^{(k)}$ and $T_{1,k}^{(k)}$ are non-zero. Therefore $\langle S_1 | S_2 \rangle > 0$, and hence $C_q \leq C_\mu$. Quantum ϵ -machines are simpler than their simplest classical alternatives, and this difference becomes ever more pronounced as the future statistics of their associated causal states becomes more similar.

II. COMPLEXITY WITH QUANTUM LOGIC

A. Quantum ϵ -machines and Complexity

How does the advent of quantum logic affect our original motivation of studying ϵ -machines? The statistical complexity was proposed as a measure of structure in line with the ideal that the more memory required to model a process the greater its complexity. The above observations indicate that if we adopt such ideals we must necessarily accept that what we perceive to be complex depends on what information theory we use.

If we are limited to classical logic then the simplest way to model a process is through classical ϵ -machines resulting in a perceived complexity of C_μ . Thus, Crutchfield defined C_μ as the statistical complexity – and motivated it as an intrinsic measure of structure and complexity of a given stochastic process $P(\overleftarrow{X}, \overrightarrow{X})$.

However, if we admit quantum information and quantum logic a stochastic process is likely to look simpler. If we are to model the process using a quantum ϵ -machine, then the perceived complexity of the system, C_q , is often strictly less than C_μ . To a creature that employs quantum ϵ -machines reality would appear simpler. Could C_q provide an alternative quantifier of complexity?

This article explores the behavior of C_q . In studying how it evolves for a simple process we show how our notions of structure and complexity can diverge in the presence of quantum logic. We outline a scenario in which the behavior of C_q aligns much more closely with our expectations of how complexity behaves.

We note, however, that while it is tempting to immediately name C_q as the quantum statistical complex-

ity, such an assignment would be rash. The question of whether quantum ϵ -machines are the simplest quantum models remains open. In this article we refer to C_q as the *quantum ϵ -machine complexity* to avoid confusion, and leave questions of whether even simpler quantum models exist for future work.

B. The Classical-Quantum Divergence

To illustrate the divergence in the complexity of classical and quantum ϵ -machines, we outline a simple process involving a box containing a single coin. At each time-step t the box is perturbed and the state of the coin is measured. This results in a binary output that is either heads ($x_t = 0$) or tails ($x_t = 1$), governed by random variable \mathbf{X}_t . The probability distribution over an infinite sequence of such measurements then defines a stochastic process $P(\overleftarrow{X}, \overrightarrow{X})$.

We assume that the act of perturbation may flip the state of the coin. The coin may be biased, such that perturbing a coin in state k will cause it to flip with probability $0 < q_k < 1$, $k \in \{0, 1\}$. Note that the special case of an unbiased coin was analyzed in [10], where it was referred to as the *perturbed coin*.

The resulting process is clearly Markovian; all information about \overleftarrow{X} is contained in X_{-1} . For $q_0, q_1 \neq 0.5$, we have two causal states. One equivalence class containing all pasts ending in “0”, signifying the last state was heads, and the other containing pasts ending in “1”. These causal states are denoted as S_0 and S_1 . The transition probabilities for the corresponding ϵ -machine are then given by $T_{jk}^{(r)} = \delta_k^r T_{jk}$, where T_{jk} are elements of the matrix

$$T = \begin{pmatrix} 1 - q_0 & q_1 \\ q_0 & 1 - q_1 \end{pmatrix}. \quad (7)$$

Together these objects define the classical ϵ -machine of the process (Fig 1).

The statistical complexity of the process is then determined by $H(\mathbf{S}) = -p_0 \log_2 p_0 - p_1 \log_2 p_1$, where $p_i = P(\epsilon(\overleftarrow{x}) \in S_i)$ for $i \in \{0, 1\}$. To find these, let




FIG. 2: A plot of both the classical statistical complexity C_μ (red, dashed) and C_q (blue, solid) for the perturbed coin, against the probability of flipping $q = q_1 = q_2$. At $q = 0.5$ the coin is completely fair and the output is random, making it unnecessary to store any information about the past to optimally predict the future. This plot illustrates the discontinuity in C_μ which jumps directly from 1 to 0 at $q=0.5$. In contrast C_q is a continuous function of q .

$\mathbf{p} = (p_0, p_1)$; then \mathbf{p} satisfies $\mathbf{p} = T\mathbf{p}$, with solution

$$\mathbf{p} = \frac{1}{q_1 + q_0} \begin{pmatrix} q_1 \\ q_0 \end{pmatrix}. \quad (8)$$

Thus $p_0 = \frac{q_1}{q_0 + q_1}$ and $p_1 = \frac{q_0}{q_0 + q_1}$. The perturbed coin in its own simplest model, with statistical complexity

$$C_\mu = -\frac{q_0}{q_0 + q_1} \log \left(\frac{q_0}{q_0 + q_1} \right) - \frac{q_1}{q_0 + q_1} \log \left(\frac{q_1}{q_0 + q_1} \right). \quad (9)$$

By comparison the quantum causal states are $|S_0\rangle = \sqrt{1 - q_0}|0\rangle + \sqrt{q_0}|1\rangle$ and $|S_1\rangle = \sqrt{q_1}|0\rangle + \sqrt{1 - q_1}|1\rangle$. Thus the complexity of the quantum ϵ -machine is $C_q = -\text{Tr } \rho \log \rho$, with $\rho = p_0|S_0\rangle\langle S_0| + p_1|S_1\rangle\langle S_1|$.

The complexity of classical ϵ -machines and their quantum counterpart are plotted in Fig 2 for the special case where the coin is unbiased, i.e., $q_0 = q_1 = q$. The classical measure of complexity is clearly discontinuous: for $q \neq 0.5$, the two causal states are equiprobable and thus $C_\mu = 1$ [6]. Yet at $q = 0.5$, the process becomes completely random and thus the future is statistical identical for all pasts; there is only one causal state and $C_\mu = 0$. An infinitesimal perturbation in q around 0.5 leads to a sudden change in statistical complexity.

The complexity of the quantum ϵ -machine, in contrast, remains continuous. As q approaches 0.5, the process becomes progressively more random; the overlap between the future statistics of the two causal states increase, thereby increasing the advantage of storing them in non-orthogonal states. In the limit $q \rightarrow 0.5$, $\langle S_0 | S_1 \rangle \rightarrow 1$ and C_q smoothly converges to 0.

Thus, what appears in ϵ -machines as a striking discontinuity in complexity vanishes when these machines are quantized. Conceptually, discontinuities in complexity appear difficult to explain. Why would an infinitesimal

perturbation in the observed statistics leads to a large change in its perceived structure and complexity? The use of quantum ϵ -machines seems to resolve this conundrum; and thus presents a promising refinement of their classical predecessors.

III. THE THERMALIZING QUBIT CLOUD - A SYSTEM OF EVER INCREASING ENTROPY

Can quantum ϵ -machines be used to formalize our intuition of how complexity should evolve? We shed light on this by studying how complexity evolves in a simple system featuring monotonically increasing entropy - a cloud of qubits undergoing gradual thermalization. Such an environment adheres to the second law of thermodynamics; mimicking our ideal that all systems graduate towards a state of total disorder.

Our intuition tells us that at either end of the scale complexity should be minimal. The system should evolve from something simple, to something more complex, and back. Should C_q be a true quantifier of complexity, we would expect it to exhibit similar behavior.

Consider an environmental bath containing a large number of ($N \gg 1$) qubits, with a global thermalization parameter λ . When $\lambda = 0$, the system is pure, and all qubits are in the state $\rho(0) = |0\rangle\langle 0|$. When $\lambda = 1$, all qubits are maximally random with $\rho(1) = I/2$. With gradual thermalization, the system, monotonically evolves from order to disorder, such that each qubit evolves according to

$$\rho_e(\lambda) = (1 - \lambda)|0\rangle\langle 0| + \frac{\lambda}{2}I. \quad (10)$$

The whole system can be described as $\rho_e^{\otimes N}(\lambda)$. Clearly this system exhibits monotonically increasing entropy, evolving smoothly from the minimum value of 0 to the maximum value of N .

To characterize how complexity evolves in this system, we introduce an observer; an entity that probes this thermalizing qubit cloud at particular values of λ . For this toy model we construct a scheme in which the observer's measurement device consists of a single probe qubit. This probe is engineered to interact with qubits within the cloud (specifics below), during which the observer monitors the probe at discrete time intervals to retrieve a sequence of binary outcomes $\cdots x_{-2}x_{-1}x_0x_1\cdots$. From these statistics the observer constructs a stochastic process $P(\vec{X}, \vec{X})$, which is then used to quantify the complexity of the environment.

The specifics of our measurement scheme are outlined in Fig. 3. At each time-step t , the observer scatters his probe qubit ρ_{obs} with a randomly selected qubit in the cloud; he then chooses one of the two output qubits with probability $0 < g < 1$ and measures it in the computational basis to generate output x_t . Repetition of this process then results in a bit-string whose distribution is




FIG. 3: Circuit representation of how the observer's qubit evolves through interactions with the qubit cloud. At each time step, t , a new ancillary environmental qubit $\rho_e(\lambda)$ interacts with the observer's probe ρ_{obs} . The interaction is modeled by a controlled unitary operation, $|0\rangle_e\langle 0|_e \otimes 1 + |1\rangle_e\langle 1|_e \otimes X\kappa$ where \otimes is the direct product and $X\kappa = \exp(i\hat{X}\kappa)$ is a function of the Pauli X operator, \hat{X} , and an interaction strength parameter $0 \leq \kappa \leq \frac{\pi}{2}$. For the maximally entangling case $\kappa = \frac{\pi}{2}$ this interaction reduces to a controlled NOT gate. The pSWAP (probabilistic SWAP) defines the execution of a SWAP gate, $U_s : |\phi\rangle|\psi\rangle \rightarrow |\psi\rangle|\phi\rangle$ with probability g ; and models the observer's uncertainty of which qubit represents the probe after interaction. The symmetric case of $g = 0.5$ corresponds to the case where the observer completely loses track of which output qubit is which. At the end of the subroutine the observer measures his qubit in the Z basis and outputs the answer as x_t . He then repeats the exercise with a new ancillary environmental qubit.

governed by an associated stochastic process $P(\overleftarrow{X}, \overrightarrow{X})$. We make two simplifying assumptions about the system.

1. The timescale in which these measurements are made is infinitesimal compared to the timescale in which λ evolves.
2. The number of qubits in the cloud, N , is sufficiently large that the probe qubit never interacts with the same qubit twice.

These assumptions ensure that at each value of λ we will probe the environment many times (collect many measurement outcomes) allowing the bit string we generate to be modeled by a stationary stochastic process. These assumptions apply, for example, in studying systems on the macroscopic scale.

Under these assumptions it is easy to see that the entropic properties of $P(\overleftarrow{X}, \overrightarrow{X})$ follow the entropic properties of the qubit cloud. At $\lambda = 0$, $P(\overleftarrow{X}, \overrightarrow{X})$ takes on unit probability for a sequence of 0's, and thus has zero entropy. Meanwhile, its entropy is maximal at $\lambda = 1$. Thus, the toy model provides a first order testing ground for our proposed measure of complexity.

A. Statistical Complexity of the Qubit Cloud

The first step in determining how complexity evolves in our thermalizing qubit cloud is to construct its associated ϵ -machine. To do this, we need to characterize the set of causal states \mathcal{S} and associated transition probabilities. The first observation we make is that the process is Markovian. At each time step the equivalent quantum circuit is initialized in input state: $|k\rangle\langle k| \otimes \rho_e(\lambda)$, where the last measurement outcome $x_{t-1} = k \in \{0, 1\}$ completely determines the observer's input state $|k\rangle$. This input determines \overrightarrow{X} , for more detail see the Appendix. Hence all information about the future \overrightarrow{X} , is contained in X_{-1} .

This Markovian property implies that the process has at most two causal states: $S_0 = \{0, 10, 00, 110, \dots\}$, the set of pasts ending with "0", and $S_1 = \{1, 01, 11, 011, \dots\}$, the set of pasts ending with "1". Indeed, provided the process is not completely random or completely uniform (i.e. $0 < \lambda < 1$), we must record something about the past to predict the future, i.e. $P(\overrightarrow{X}|S = S_0) \neq P(\overrightarrow{X}|S = S_1)$. Thus the two causal states S_0 and S_1 are distinct.

To evaluate the transition probabilities $T_{jk}^{(r)}$, we first note that $T_{jk}^{(r)} = \delta_k^r T_{jk}$: if the system transits to S_k at time t it will always emit $x_t = k$. Thus the thermalizing qubit cloud generates a process $P(\overleftarrow{X}, \overrightarrow{X})$ that is statistically identical to the perturbed coin (see Fig. 1) with T_{jk} described by the transition matrix

$$T = \begin{pmatrix} 1 - q_0 & q_1 \\ q_0 & 1 - q_1 \end{pmatrix}, \quad (11)$$

where q_0 and q_1 are λ -dependent. We evaluate these in the Appendix A to find

$$q_0(\lambda) = g \frac{\lambda}{2} + (1 - g) \frac{\lambda}{2} \sin^2(\kappa), \quad (12a)$$

$$q_1(\lambda) = (1 - g) \frac{\lambda}{2} \sin^2(\kappa) + g \left(1 - \frac{\lambda}{2}\right). \quad (12b)$$

The statistical complexity of the process is then defined by the complexity of this model; i.e., the entropy of the random variable \mathbf{S} over causal states. Let $p_i = P(\mathbf{S} = S_i)$ be the probability an ϵ -machine is in state S_i , then from Eq. (8);

$$p_0(\lambda) = \frac{q_1(\lambda)}{q_0(\lambda) + q_1(\lambda)}, \quad p_1(\lambda) = \frac{q_0(\lambda)}{q_0(\lambda) + q_1(\lambda)}, \quad (13)$$

and the statistical complexity is thus

$$C_\mu(\lambda) = - \sum_i p_i(\lambda) \log_2 p_i(\lambda), \quad 0 < \lambda < 1. \quad (14)$$

for the special case where $\lambda = 1$, $P(\overleftarrow{X}, \overrightarrow{X})$ becomes completely random and there is only a single causal state encompassing all possible pasts. Therefore the complexity of the system reduces to 0, i.e., $C_\mu(1) = 0$.

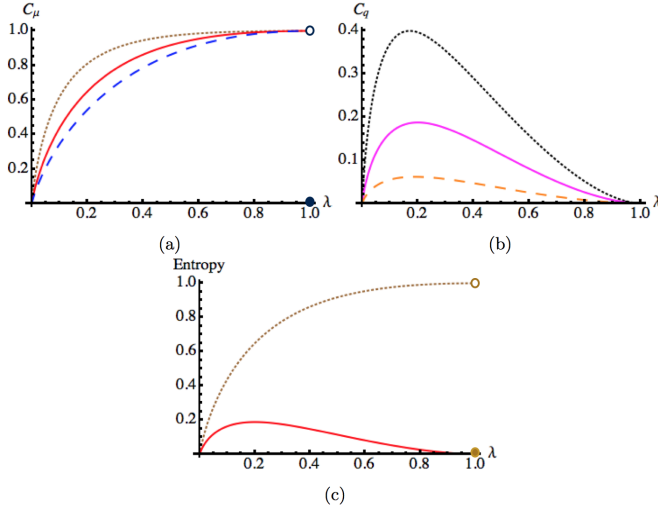


FIG. 4: Overview of C_q and C_μ for the thermalizing qubit cloud, as a function of the thermalization parameter, $0 < \lambda < 1$. As λ increases more noise is introduced and the environmental qubits, $\rho_e(\lambda)$, becomes more mixed. (a) Plot of C_μ vs. λ for various swap probabilities, $g = 0.25$ (brown, dotted), $g = 0.5$ (red, solid) and $g = 0.75$ (blue, dashed). Qualitatively C_μ is a monotonically increasing function of λ . (b) Plot of C_q vs. λ for values of $g = 0.25$ (black, dotted), $g = 0.5$ (purple, solid) and $g = 0.75$ (orange, dashed). C_q dies down as $\lambda \rightarrow 1$ and the corresponding stochastic process becomes more random. (c) Plot of $g = 0.5$ case for both C_μ (brown, dotted) and C_q (red, solid) to illustrate the comparative behaviour.

B. Quantum ϵ -machines and their Complexity

The quantum ϵ -machine can be determined by directly quantizing the causal states; resulting in two quantum causal states

$$|S_0\rangle = \sqrt{1 - q_0}|0\rangle + \sqrt{q_0}|1\rangle, \quad (15a)$$

$$|S_1\rangle = \sqrt{q_1}|0\rangle + \sqrt{1 - q_1}|1\rangle. \quad (15b)$$

The complexity of the quantum ϵ -machine is given by the von Neumann entropy of the density operator

$$\rho = p_0|S_0\rangle\langle S_0| + p_1|S_1\rangle\langle S_1|, \quad (16)$$

where p_i is the probability of finding the quantum ϵ -machine in the state $|S_i\rangle$. This probability satisfies Eq. (13).

C. Complexity Dynamics

We now have all the tools necessary to analyze how complexity evolves within the qubit cloud—both through the lens of conventional ϵ -machines and their quantum mechanical counterpart.

Firstly, observe that the complexity of the system, regardless of which measure we use (C_μ or C_q), is indeed

minimal for $\lambda = 0, 1$. Eq. (12) implies that when $\lambda = 0$, $q_0(0) = 0$ and $q_1(0) > 0$; thus the ϵ -machine resides in state S_0 with unit probability. The process is trivial (the output is a uniform string of 0s), resulting in $C_\mu = C_q = 0$. On the other end, $q_0 = q_1$ when $\lambda = 1$, thus $P(\vec{X}|\mathbf{S} = S_0) = P(\vec{X}|\mathbf{S} = S_1)$. This indicates that the two causal states have statistically identical futures; and therefore collapse into a single causal state that encompasses all pasts. Again $C_\mu = C_q = 0$. Thus, both complexity measures agree with our intuition that a system at either 0 or maximal entropy process is completely trivial.

The behavior of C_q and C_μ , however, diverges for intermediate values of λ . Fig. 4 displays the generic behavior of both C_μ and C_q . C_μ is a monotonically increasing function for all $\lambda \neq 1$. In fact, $\lambda = 1$ is a point of discontinuity; where C_μ drops sharply to 0. It indicates that the process is ‘maximally complex’ an infinitesimal distance away from being completely random. If we put this in the context of a physical system where complete randomness is an idealized limit, it seems to suggest that complexity shares almost identical behavior with entropy. These results contrast quite sharply with our perceived conception of complexity, and intuition that complexity should be complementary to randomness [9, 11]. The behavior of C_q , on the other hand, appears in line with what we expect. The quantum ϵ -machine complexity rises and falls continuously, taking on a peak value at $\lambda \approx 0.2$. There is no sudden jumps, and the complexity of the qubit cloud near complete thermalization is indeed very close to 0.

Fig. 4 (c) compares C_q and C_μ directly for the special case where $g = \frac{1}{2}$ and $\kappa = \pi/2$. This corresponds to the limiting scenario where the probe interacts with the qubit cloud through the idealized CNOT gate and the observer measures one of the two outputs at random. It is clear that the quantum ϵ -machine is able to model the resulting process using far less memory; and this advantage grows with λ . These results are not isolated to specific parameter choices. Fig. 5 displays the behavior of C_q for various values of κ and g . In general, C_q is a continuous function of λ and gradually diminishes as λ approaches either 0 or 1, while peaking somewhere in the middle.

These results can be understood qualitatively. As the process becomes more random there are two contributing sources to the complexity. The first is positive: noise gives the system the energy to transition from causal state S_0 to S_1 . This applies an equalizing pressure to the population level of two causal states S_0 , and S_1 . Thus, more memory is required to distinguish the two possibilities. The second a negative contribution. As noise increases; the future statistics of S_0 and S_1 become less distinct, and thus distinguishing between the two is less meaningful. The evolution of complexity within the system rests on the balance of these two contributions.

The classical statistical complexity, however, accounts only for the former contribution. It is based entirely on the entropy of distinguishing causal states S_0 and S_1 .

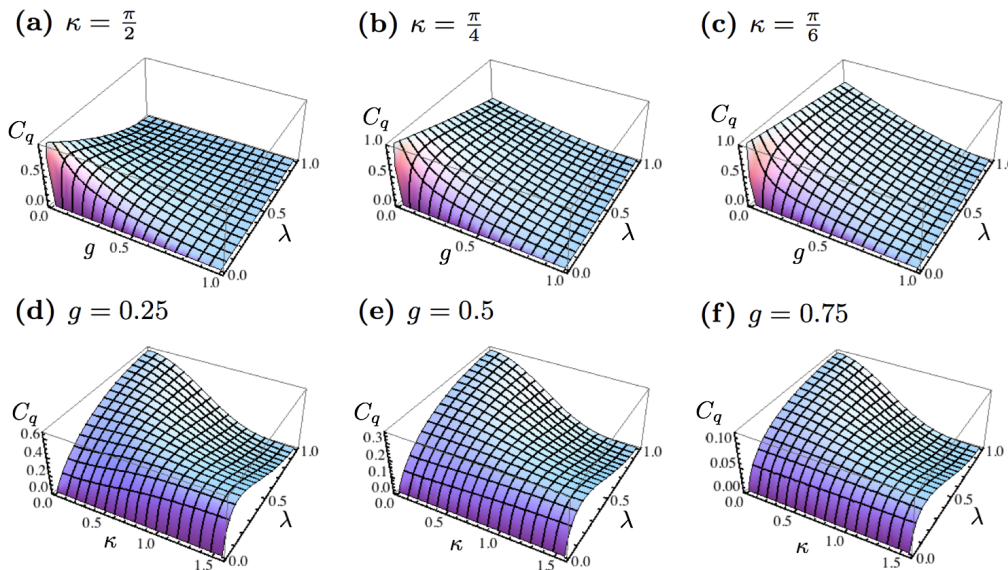


FIG. 5: Evolution of C_q with thermalization parameter λ for various values of interaction strength κ and swap probability g . In (a), (b), (c), g is varied for various fixed values of κ . In (d), (e), (f), κ is varied for various fixed values of g .

Within our toy model as λ increases $p_0 \rightarrow p_1$. Thus the statistical complexity undergoes monotonic gain. The progressive convergence in the conditional future of these two causal states remains ignored until $\lambda = 1$. C_q , on the other hand, continually accounts for the latter contribution. As the conditional futures of S_0 and S_1 converge, so do the quantum causal states $|S_0\rangle$ and $|S_1\rangle$. Thus the amount of memory required to model the process evolves continuously reflecting our ideal that complexity rests on the balance between order and disorder.

IV. DISCUSSION

In this article we explored a quantifier of complexity by extending the framework of ϵ -machines into the quantum mechanical regime. In computational mechanics the information content of ϵ -machines presents a popular approach to quantifying the structure of a given stochastic process, the rationale being that they are its simplest models. The advent of quantum ϵ -machines, however, demonstrated that simpler models do exist. This motivated us to ask “how complex would a stochastic process look to a quantum ϵ -machine?”

Our results demonstrated a marked divergence in the complexity of ϵ -machines and their quantum mechanical counterparts. In a process at various different stages of thermalization λ we found that the statistical complexity increased monotonically with λ . Only at infinite temperature did it drop discontinuously to zero. The quantum ϵ -machine complexity, on the other hand, behaved as a smooth function of λ first rising for low values of entropy,

then falling as the system approached total randomness.

When we envision the dynamics of complexity the latter quantity seems more reasonable. If we are to estimate the probability distribution of a given stochastic process $P(\vec{X}, \vec{X})$ through observations, there will always be a statistical margin of error. Thus, it is unsatisfactory for a quantity that describes the process complexity to jump to a fixed value (here 1) for an infinitesimal perturbations away from total randomness. The continuity in the complexity of quantum ϵ -machines is thus a welcome trait.

This article only skims the surface of how quantum theory may combine with computational mechanics. Many open questions remain. On the one hand, the model that we have presented here is but a toy with very specific assumptions on how the qubit cloud was probed. Could similar techniques be applied to more complex systems that have been studied within the framework of computational mechanics? We note that previous studies of statistical complexity in the Ising lattice demonstrated similar behavior [6]. The complexity rose monotonically with temperature and only dropped to 0 via a sudden discontinuous jump at infinite temperature. Could quantizing ϵ -machines remove this discontinuity?

On the other hand there is currently no proof that quantum ϵ -machines are the provably simplest quantum models. If this turns out false, would the true minimal amount of memory, C_Q required to simulate a given stochastic process share the qualitative features of C_q ? Certainly, using bounding arguments (As $0 \leq C_Q \leq C_q$; we can see that C_Q has to be continuous at the point of maximal randomness in the qubit cloud.

The ultimate goal would be to present an operationally

meaningful, yet nevertheless computable, quantifier of complexity. This would substantiate our intuition that complexity lies at the border between order and chaos - and thus pave the path for developing universal laws that govern complexity. This article presents one clue in the big puzzle - our notions of what is complex is affected by what sort of information we use; and quantum information could be a valuable tool in understanding what around us is ultimately complex.

Acknowledgements.— The authors acknowledges helpful discussions with Karoline Weisner, Alex Monras and Borivoje Dakić. This work is supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00302, the National Natural Science Foundation of China Grant 61033001, 61361136003, the Singapore Ministry of Education and the Academic Research Fund Tier 3 MOE2012-T3-1-009 “Random numbers from quantum processes”.

-
- [1] A. N. Kolmogorov, *Theoretical Computer Science* **207**, 387 (1998).
 - [2] J. Ladyman, J. Lambert, and K. Wiesner, *European Journal for Philosophy of Science* **3**, 33 (2013).
 - [3] J. P. Crutchfield and K. Young, *Physical Review Letters* **63**, 105 (1989).
 - [4] C. R. Shalizi and J. P. Crutchfield, *Journal of statistical physics* **104**, 817 (2001).
 - [5] C.-B. Li, H. Yang, and T. Komatsuzaki, *Proceedings of the National Academy of Sciences* **105**, 536 (2008).
 - [6] J. P. Crutchfield and D. P. Feldman, [arXiv:9702191] (1997).
 - [7] K. Wiesner, M. Gu, E. Rieper, and V. Vedral, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science* **468**, 4058 (2012).
 - [8] W. G. Rory Cerbus, [arXiv:1403.5356] (2014).
 - [9] J. P. Crutchfield, *Nature Physics* **8**, 17 (2012).
 - [10] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, *Nature communications* **3**, 762 (2012).
 - [11] J. P. Crutchfield, *Physica D: Nonlinear Phenomena* **75**, 11 (1994).
 - [12] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).

Appendix

A. Evaluation of Perturbation Parameters for the thermalizing qubit cloud

In this Appendix we give a more in depth treatment of the thermalizing qubit cloud. The circuit representation of this model can be broken down into two stages: a two-qubit interaction between the observer’s qubit and the environmental qubit, followed by a probabilistic swap defined in the caption of Fig. 3. Formally during the first stage the two-qubit unitary is given by

$$C_{X\kappa} = 1 \otimes |0\rangle_e \langle 0|_e + X\kappa \otimes |1\rangle_e \langle 1|_e, \quad (17)$$

where $X\kappa = \exp(i\hat{X}\kappa)$ is defined in terms of the Pauli X operator, and \otimes denotes the direct product.

The probabilistic SWAP operation acts on the combined system consisting of the environmental qubit and the observer’s qubit, which we denote by $\rho_{obs,e}$. We define this transformation through

$$\rho_{obs,e} \rightarrow \rho'_{obs,e} = gU_S\rho_{obs,e}U_S^\dagger + (1-g)\rho_{obs,e} \quad (18)$$

where U_S is the standard unitary SWAP operation defined by $U_S|\phi\rangle_{obs}|\psi\rangle_e = |\psi\rangle_{obs}|\phi\rangle_e$, and g parametrises the probability of swapping.

The assumption that the bath is extremely large, such that the observer’s qubit never interacts with the same environmental qubit twice, implies that the system is initialized in a product state $\rho_{obs} \otimes \rho_e(\lambda)$. At each time step, t , the observer’s qubit is initialized in the quantum state corresponding to the last measurement outcome x_{t-1} , while the environmental qubit is given by (10).

Explicitly, if the outcome of the last measurement is $x_{t-1} = k$ (for $k \in \{0, 1\}$), then the observer’s qubit is initialized in state $|k\rangle$ and the combined two-qubit system, $\rho_{obs,e}$, is initialized as

$$\rho_{obs} \otimes \rho_e(\lambda) = |k\rangle\langle k| \otimes \left[(1-\lambda)|0\rangle\langle 0| + \frac{\lambda}{2}\mathbb{1} \right] \quad (19)$$

$$= (1 - \frac{\lambda}{2})|k0\rangle\langle k0| + \frac{\lambda}{2}|k1\rangle\langle k1|. \quad (20)$$

If $x_{t-1} = 0$ then after going through the $C_{X\kappa}$ interaction and the probabilistic SWAP operation we take the partial trace over the environmental qubit to recover the state of observer’s qubit directly before measurement:

$$\begin{aligned} \rho'_{obs} &= \text{Tr}_e(\rho'_{obs,e}) = \left[\left(1 - \frac{\lambda}{2}\right) + \frac{\lambda}{2} \cos^2 \kappa \right] |0\rangle\langle 0| + \\ &\left[g\frac{\lambda}{2} + (1-g)\frac{\lambda}{2} \sin^2 \kappa \right] |1\rangle\langle 1| + i \sin \kappa \cos \kappa (|1\rangle\langle 0| - |0\rangle\langle 1|). \end{aligned} \quad (21)$$

Correspondingly, if $x_{t-1} = 1$ then

$$\begin{aligned} \rho'_{obs} &= \text{Tr}_e(\rho'_{obs,e} \otimes \rho'_e) \\ &= \left[g\left(1 - \frac{\lambda}{2}\right) + (1-g)\frac{\lambda}{2} \sin^2 \kappa \right] |0\rangle\langle 0| + \\ &\left[g\frac{\lambda}{2} + (1-g)\left(1 - \frac{\lambda}{2} \sin^2 \kappa\right) \right] |1\rangle\langle 1| + \\ &(1-g)i \sin \kappa \cos \kappa (|0\rangle\langle 1| - |1\rangle\langle 0|). \end{aligned} \quad (22)$$

The terms before $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ give the probabilities of measuring $|0\rangle$ and $|1\rangle$ and hence the statistics of the next output. The Markovian nature of the ϵ -machine is clearly demonstrated by the fact that the probability $x_t = 0$ (or 1) depends only on the value of x_{t-1} , the last

measurement outcome. This establishes the set of causal states as $S_0 = \{0, 10, 00, 110, \dots\}$, the set of pasts ending with “0”, and $S_1 = \{1, 01, 11, 011, \dots\}$, the set of pasts ending with “1”. From these results we can also find the transition probabilities for the corresponding stochastic process.

B. The Special Case of Maximally Interacting Probes

It is instructive to first outline special case where where $\kappa = \pi/2$. Using the value of $\kappa = \pi/2$ in Eq. (21) and (22) we can simplify the transition probabilities to:

$$P(\mathbf{S}_t = S_0 | \mathbf{S}_{t-1} = S_0) = 1 - \frac{\lambda}{2}, \quad (23a)$$

$$P(\mathbf{S}_t = S_1 | \mathbf{S}_{t-1} = S_0) = \frac{\lambda}{2}, \quad (23b)$$

$$P(\mathbf{S}_t = S_0 | \mathbf{S}_{t-1} = S_1) = g \left(1 - \frac{\lambda}{2}\right) + (1 - g) \frac{\lambda}{2}, \quad (23c)$$

$$P(\mathbf{S}_t = S_1 | \mathbf{S}_{t-1} = S_1) = g \frac{\lambda}{2} + (1 - g) \left(1 - \frac{\lambda}{2}\right). \quad (23d)$$

The ϵ -machine of the process is presented in Fig. 6. The




FIG. 6: Causal state diagram for the thermalizing qubit cloud with interaction strength set by $\kappa = \frac{\pi}{2}$. The two causal states are denoted S_0 and S_1 and an arrow from S_j to S_k represents the corresponding transition, with label denoting the transition probability T_{jk}^r and corresponding emission r .

quantum causal states are subsequently reduced to

$$|S_0\rangle = \sqrt{1 - \frac{\lambda}{2}} |\widetilde{0}\rangle + \sqrt{\frac{\lambda}{2}} |\widetilde{1}\rangle, \quad (24)$$

$$|S_1\rangle = \sqrt{g(1 - \frac{\lambda}{2}) + (1 - g) \frac{\lambda}{2}} |\widetilde{0}\rangle + \sqrt{g(\frac{\lambda}{2}) + (1 - g)(1 - \frac{\lambda}{2})} |\widetilde{1}\rangle. \quad (25)$$

The resulting state of the quantum ϵ -machine is:

$$\rho = p_0 |S_0\rangle \langle S_0| + p_1 |S_1\rangle \langle S_1|. \quad (26)$$

Substituting $q_0 = \frac{\lambda}{2}$ and $q_1 = g(1 - \frac{\lambda}{2}) + (1 - g) \frac{\lambda}{2}$ into Eq. (8) directly yields

$$p_0 = \frac{-2g(\lambda - 1) + \lambda}{2(g + \lambda - g\lambda)}, \quad (27a)$$

$$p_1 = \frac{\lambda}{2(g + \lambda - g\lambda)}. \quad (27b)$$